Software Defined Data Protection mit  IBM **Spectrum Protect**  IBM **Spectrum Scale**

Ende-zu-Ende Datensicherungs-Architektur bei einem Pharmaunternehmen

Flexibler Speicher für Spectrum Protect

08. März 2017

Stefan Schädeli, *SME Storage and Backup*
Andre Gaschler (gaschler@de.ibm.com), *Consulting IT Specialist Data Protection & Retention Solutions*

The environment we will speak about

- Global operating pharmaceutical company
  - 100k+ employees
  - 200+ IT relevant locations
- IT overview
  - Global strategic datacenters
    - In pairs serving regions
  - Division IT Organization in integration

Overview old environment (without Tape Libraries)

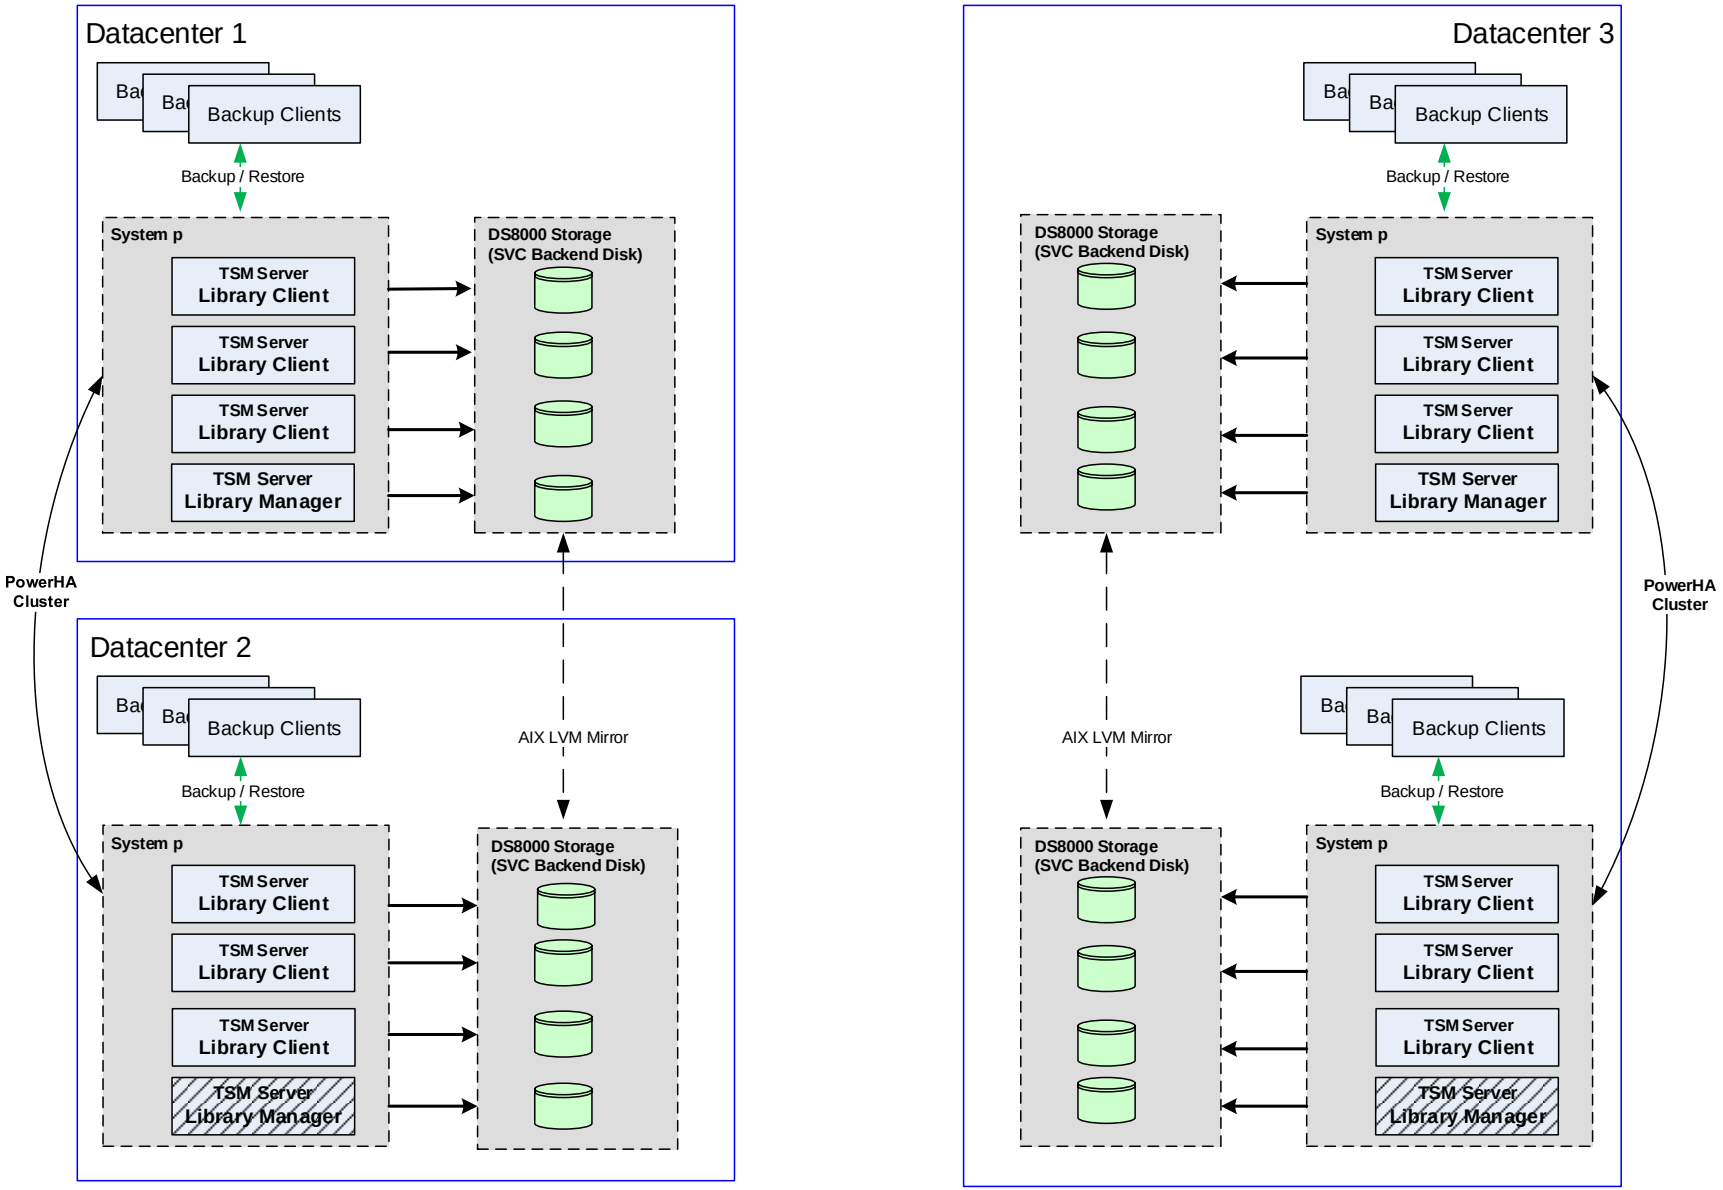

Drivers and pain points for new environment

- Existing environment was outdated on hardware and software
- Existing Resources unbalanced and overloaded
  - Either to small or to large disk storage pools
  - Not utilizing tape performance on its best when migrating data from disk to tape
  - Not utilizing LANfree for larger Databases Backups
  - Issues to perform all backups during the backup window
- Inflexible System Management
  - Continues SAN storage changes to full fill changed requirements for the different TSM servers
  - Backup data distributed over 5 Tape Libraries not following a placement policy
- No integration of smaller locations into the global solution
- Standardization is missing, no clear policies
  - Backup volume growth year-over-year of more as 1 PB

What was our vision

- Easy to manage (in operations)
  - Standardized, flexible and global backup design based on building blocks
  - Achieve transparency over backup consumers
- Flexibility
  - Seamless scalability of storage capacity and performance
  - Better storage utilization – multiple Spectrum Protect server share the same storage
  - Independently store and transparently move data: private, hybrid or public cloud
- Leverage the resources we have
  - Backup SAN with backup storage
  - Tape
- Spend wisely

Spectrum Protect with and without Spectrum Scale

Without Spectrum Scale

- Each backup server has its dedicated LUNs
- Each backup server has its own isolated file system
- Storage islands appear with underutilized capacity
- Capacity and performance management is challenging
- Scaling and performance may impact TSM

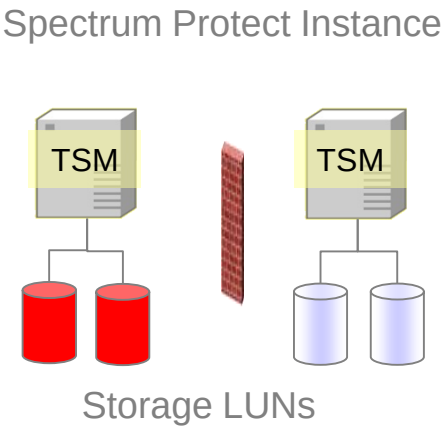

With Spectrum Scale

- Backup servers share LUNs and file systems
- Scale capacity and performance seamlessly and transparently to TSM under the shared file system global namespace
- File system replication is included
- Central administration of all storage

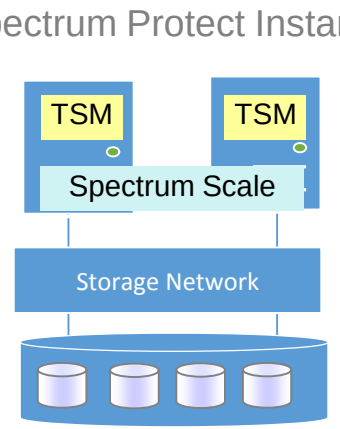

Spectrum Protect on Spectrum Scale - Overview

- Multiple Spectrum Protect (TSM) instances store DB and storage pools in a Spectrum Scale file system (GPFS)
  - Spectrum Scale provides global name space for all Spectrum Protect instances
  - Instances share all file system resources
- Spectrum Protect instances run on cluster nodes accessing the file system and disk directly
- Spectrum Scale file systems balances the workload and capacity for all TSM instances on disk
- Provides standardized, scalable and easy to use storage infrastructure for the multiple instances

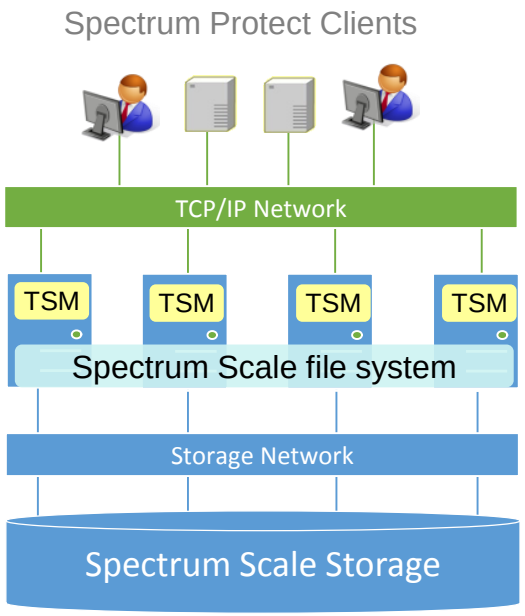

Spectrum Scale storage for Spectrum Protect

Architectural Overview Diagram of the Spectrum Protect Environment

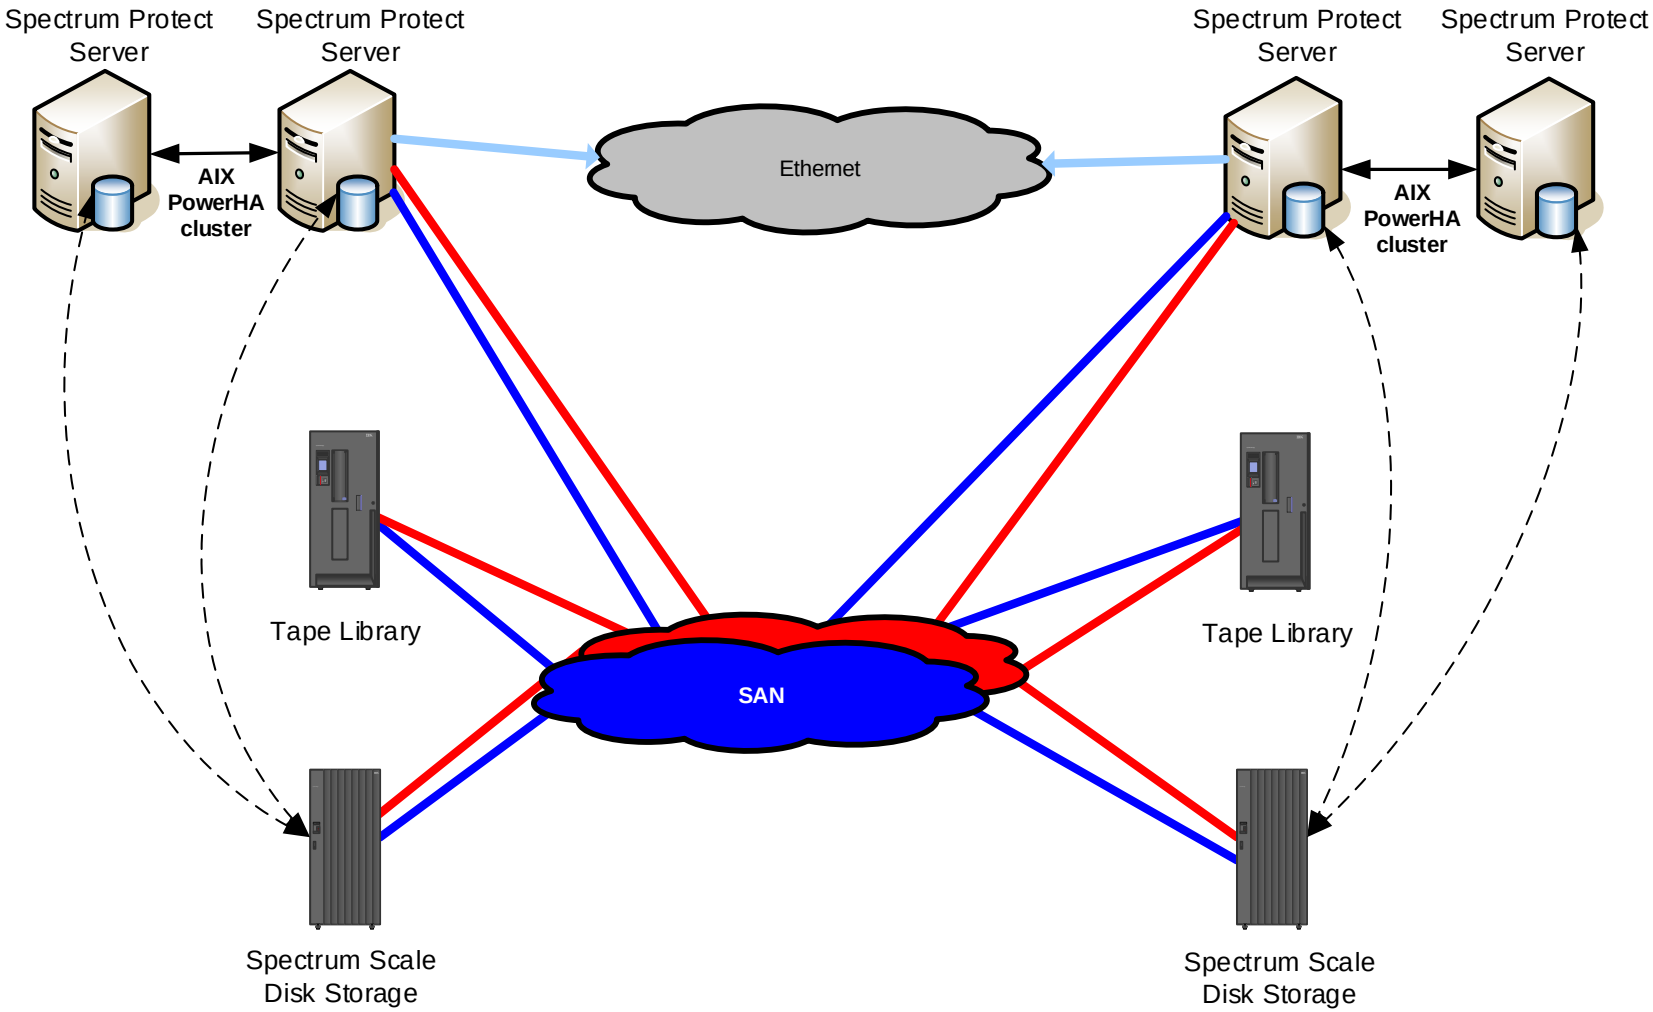

What was done to enhance the current Spectrum Protect Environment?

- Replace the existing 4 POWER7 servers with 4 new POWER8 servers
  - Install and update current servers to the latest Spectrum Protect release V7.1.x on the new HW
  - Deploy Spectrum Protect Operation Center for centralized management of the Backup environment
- Integrate Spectrum Scale (GPFS) as storage file system used by Spectrum Protect
  - Life cycle of the SVC to DH8 (introduction of compression)
  - Replace the existing DS8k storage system with new Flash System 900
- Use two IBM Enterprise Tape Libraries (ATL) for the Shared Spectrum Protect environment
  - Consolidate to 1 library per location
  - Refresh tape drives from TS1130 to TS1140 technology
  - Add the Library High Availability feature (2nd Gripper)
  - Reuse Standard and High Density (HD)

Spectrum Protect Shared Environment

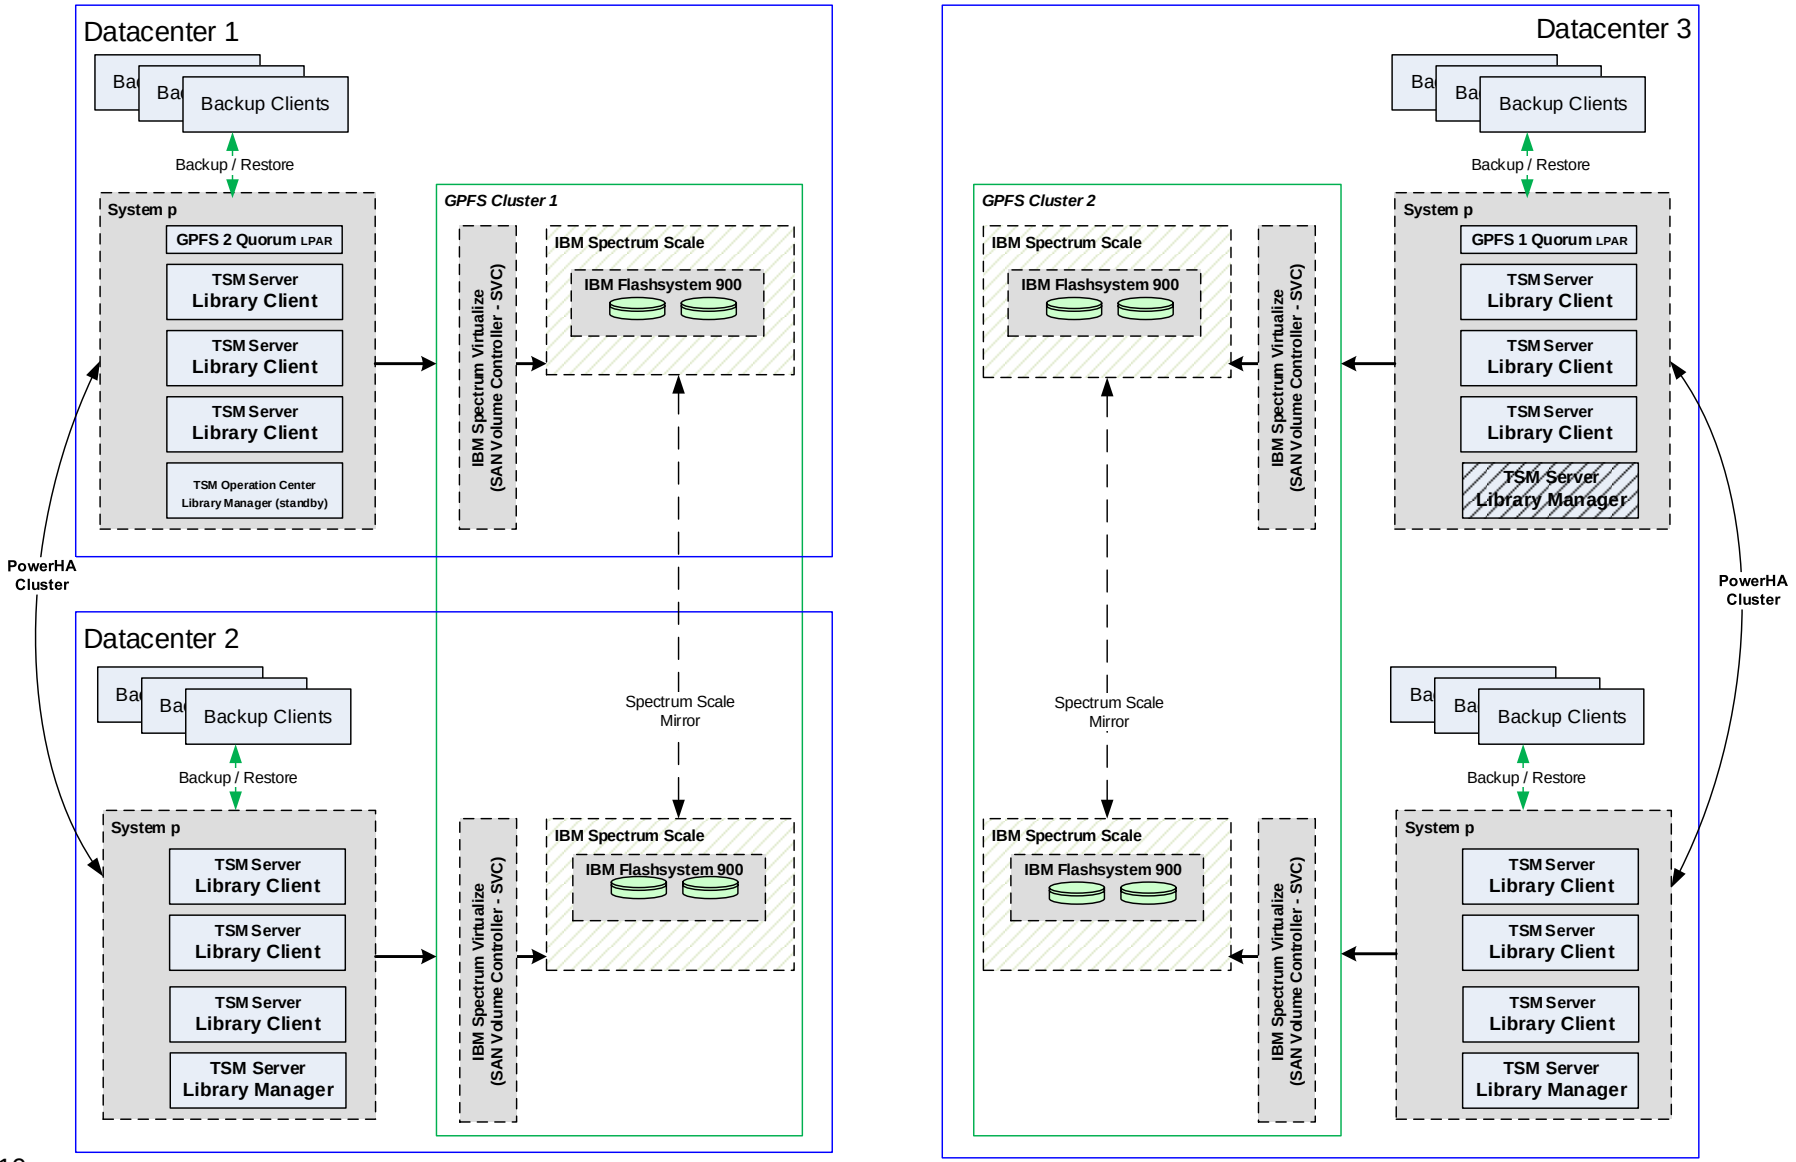

Spectrum Protect Shared Environment – Detailed View with Server Roles

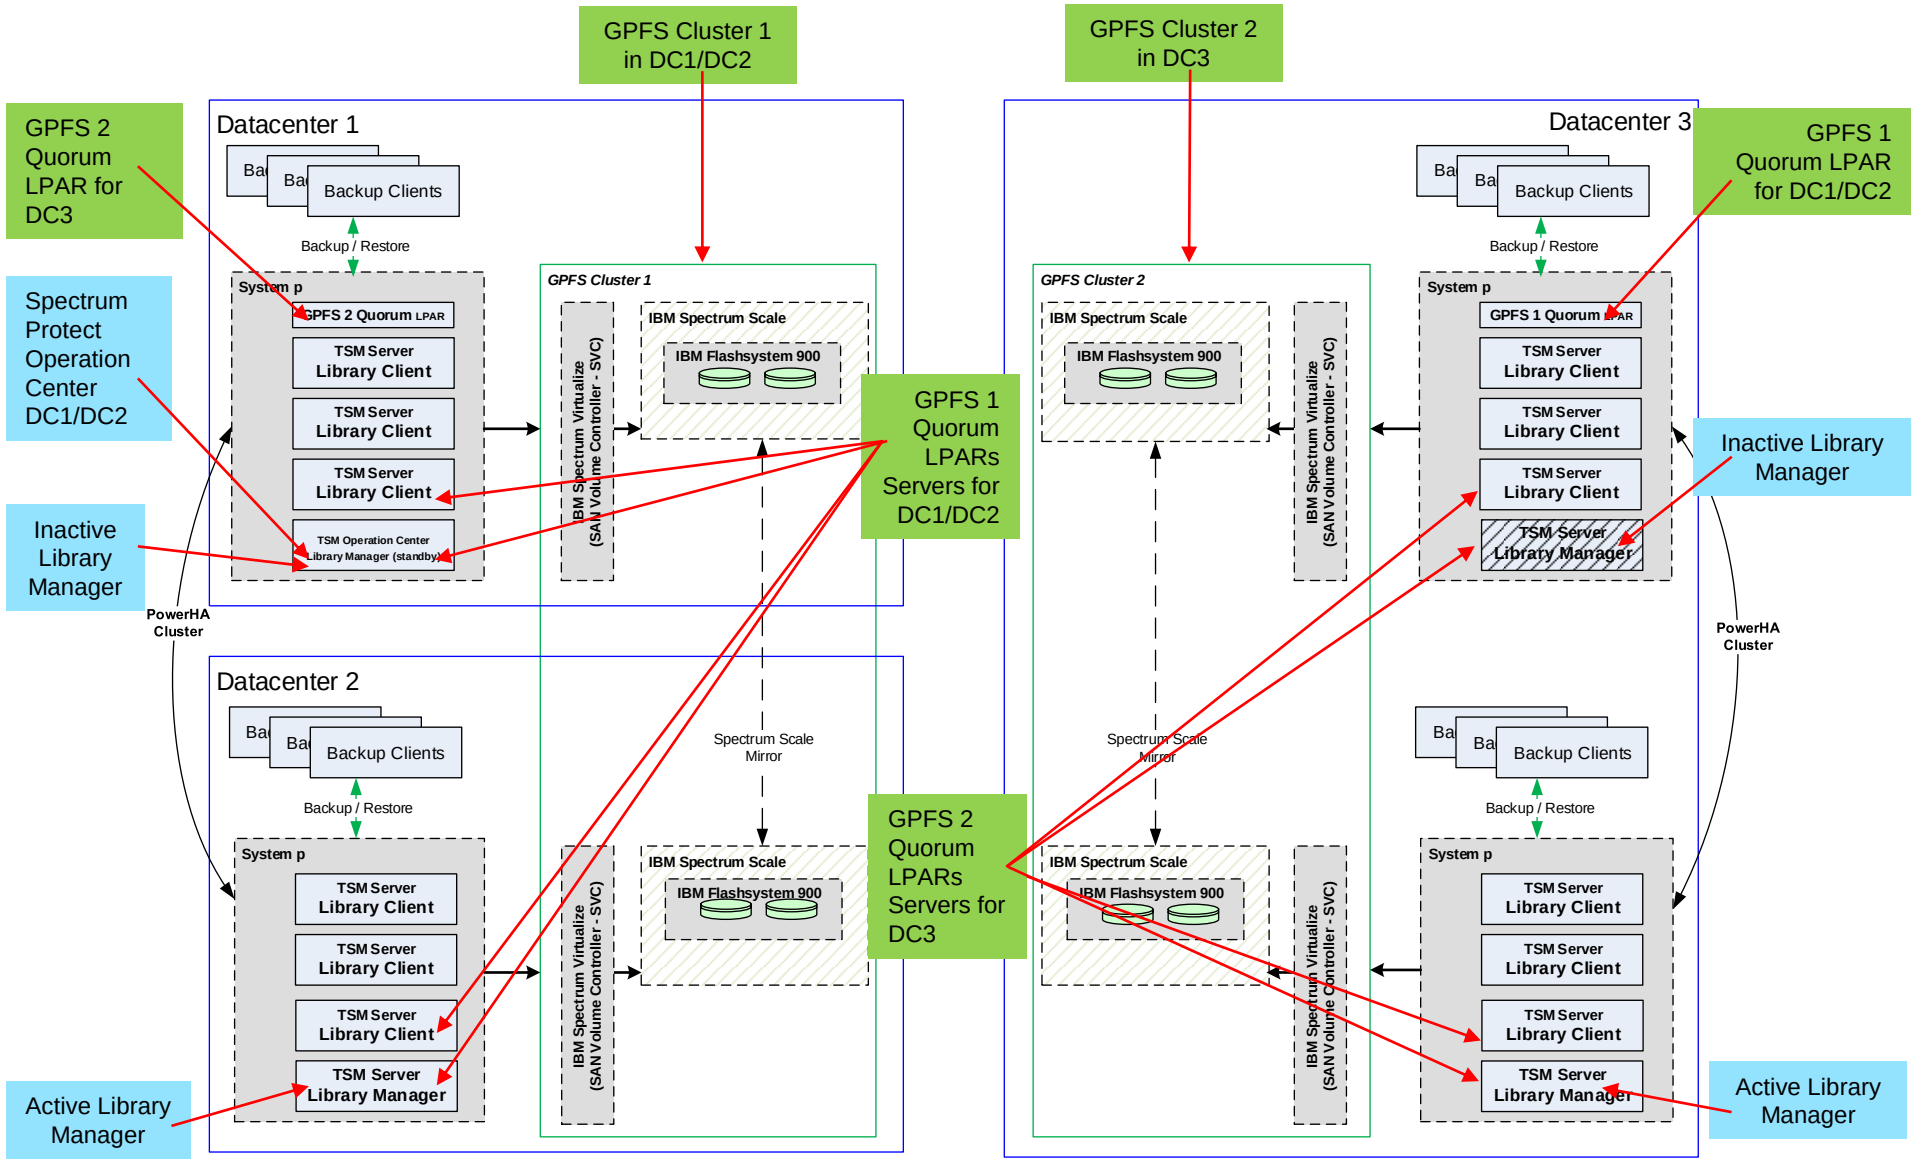

Spectrum Scale (GPFS) Design Details

- Two GPFS cluster are build (one in DC1/DC2, one in DC3) using the SVC/Flash storage, each Spectrum Protect LPAR is configured as NSD server
- Five Spectrum Scale Quorum Nodes deployed per Cluster – Due to operational reasons this was changed to 9 Quorum Nodes
  - One dedicated LPAR in DC1/ DC2 is used as Remote-Quorum for DC 3 (client-quorum)
  - One dedicated LPAR in DC3 is used as Remote-Quorum for DC1/DC2 (client-quorum)
- Four Spectrum Scale Manager Nodes deployed – Due to operations reasons this was changed to 8 Manager Nodes
- Each Spectrum Scale cluster does have 4 GPFS filesystems (DB, Active Log, Archive Log & STG)
  - *tsmdb* for all TSM databases
  - *tsmactlog* for all TSM active logs
  - *tsmarchlog* for all TSM archive logs
  - *tsmstg* for all TSM storage pools, instance directory and archive failover logs

Spectrum Protect Shared Environment – Future Extension

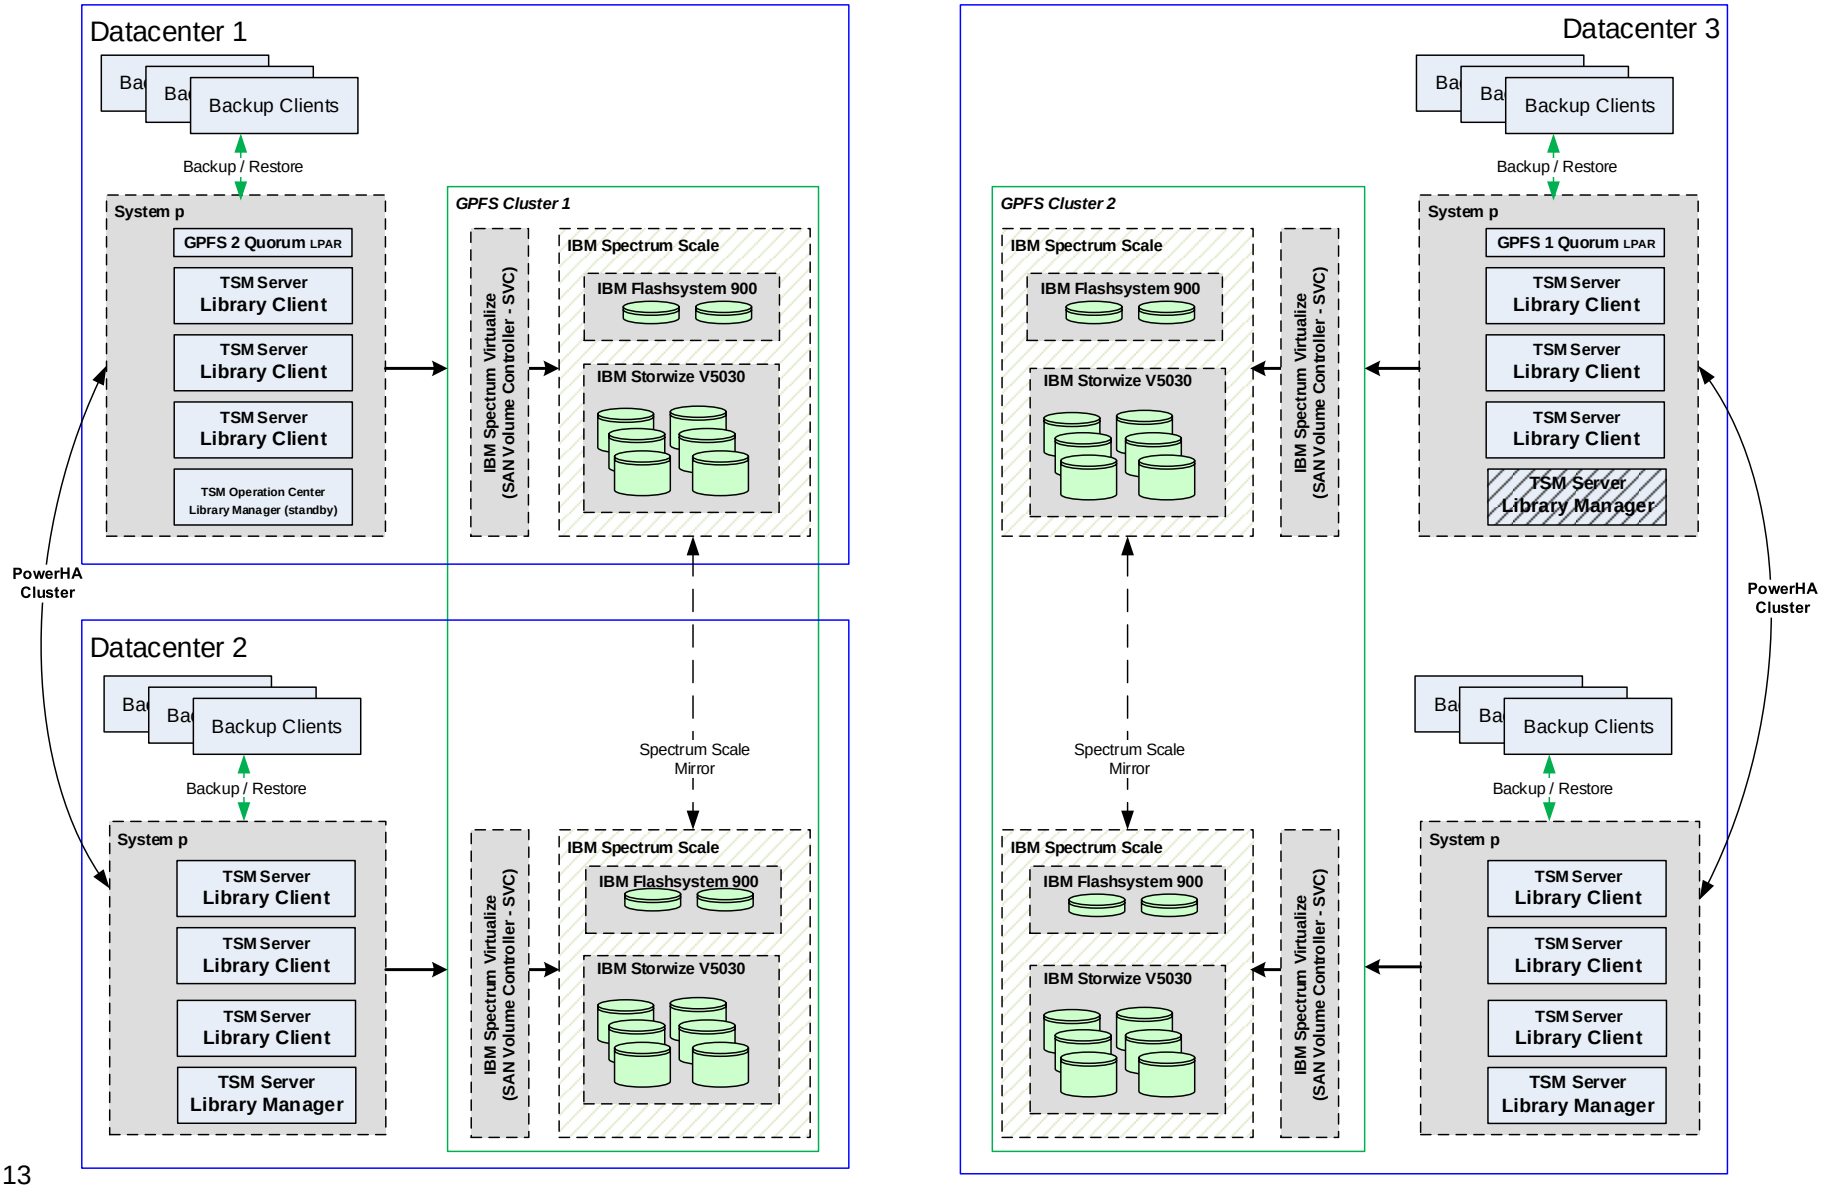


Thank You

Additional References

IBM Techdocs Library

- IBM Spectrum Protect with IBM Spectrum Scale - Introduction

- Scale out backup with TSM and GSS: Performance test results
- Elastic Storage with GPFS Native RAID performance test results with Tivoli Storage Manager over 40 GBit Ethernet
- Is a scale out NAS system the same as a scale out file system?
- Software Defined Data Protection – The next generation of backup
- LAN-free and Server-free backup to disk with IBM Spectrum Protect and IBM Spectrum Scale

The Register®

- Mirror, mirror on the wall, who has the best TSM backend of all?
- Big Blue stuffs data into backup at GIGABYTES/sec